

Wayfinding Strategies and Behaviors in Large Virtual Worlds

Rudolph P. Darken

Naval Research Laboratory
Washington, D.C. 20375
+1-202-767-2236
darken@enews.nrl.navy.mil

John L. Sibert

The George Washington University
Washington, D.C. 20052
+1-202-994-4953
sibert@seas.gwu.edu

ABSTRACT

People have severe problems wayfinding in large virtual worlds. However, current implementations of virtual worlds provide little support for effective wayfinding. We assert that knowledge about human wayfinding in the physical world can be applied to construct aids for wayfinding in virtual worlds. An experiment was conducted to determine whether people use physical world wayfinding strategies in large virtual worlds. The study measures subject performance on a complex searching task in a number of virtual worlds with differing environmental cues. The results show that subjects in the treatment without any additional cues were often disoriented and had extreme difficulty completing the task. In general, subjects' wayfinding strategies and behaviors were strongly influenced by the environmental cues in ways suggested by the underlying design principles.

Keywords

Virtual worlds, wayfinding, navigation, environmental design, spatial orientation, cognitive maps

INTRODUCTION

Problems associated with wayfinding have been encountered in every large virtual world. These problems may manifest themselves in a number of ways. Virtual world navigators may wander aimlessly when attempting to find a place for the first time. They may then have difficulty relocating places recently visited. They are often unable to grasp the overall topological structure of the space. Any time an environment encompasses more space than can be viewed from a single vantage point, these problems will occur. Compounding the problem is the fact that many virtual worlds are rarely, if ever, revisited offering the navigator no opportunity to develop a usable cognitive map of the environment.

As important as these issues are to human performance in virtual environments and the eventual acceptance of virtual environment solutions to real problems, currently, only ad hoc remedies and case studies are available for guidance.

The overall objective of our research program is to develop design principles which can be incorporated into a methodology for the design of wayfinding augmentations to virtual worlds and which will facilitate expert-like navigation performance in novice users. This paper presents an analysis of searching behaviors observed in an experiment applying real-world wayfinding and environmental design principles

to the design of virtual worlds. Our intent is to show that these same principles will produce skilled wayfinding behavior* in virtual worlds.

BACKGROUND

The principles we are investigating are based on spatial knowledge theory and environmental design methodology. Our intent is not to find new principles which are unique to virtual worlds but rather to determine how much of what is already known about wayfinding in the physical world is independent of the type of space and therefore can be applied to abstract computer-generated environments.

Spatial Knowledge

Wayfinding tasks in general require that the navigator be able to conceptualize the space as a whole. This type of configurational or topological knowledge is defined by Thorndyke [10] as *survey knowledge*. Object locations and inter-object distances are encoded in terms of a geocentric, fixed, frame of reference. This is significantly different from *procedural knowledge* which is defined as the sequence of actions required to follow a particular route. The route may make use of *landmark knowledge* which is static information about the visual details of a specific location.

Survey knowledge is map-like in nature. Accordingly, it can be acquired directly from map use. However, survey knowledge acquired from a map tends to be orientation-specific. In contrast, prolonged exposure to navigating an environment directly is more likely to result in survey knowledge which is orientation-independent.

The relative inflexibility of spatial knowledge acquired only from maps led Levine [6] to study the effect of this phenomenon on map design. He found that in order to facilitate efficient map use, the map must be congruent with the environment it represents. This is illustrated in the *forward-up equivalence* principle which states that the upward direction on a map (assuming it is perpendicular to the floor) must always show what is in front of the viewer.

Survey knowledge is hierarchical in nature [9]. Rather than encode the absolute positions and directions to every place encountered, fewer large, general, logically selected places (e.g. Washington, D.C.) are encoded with subnetworks of smaller, more specific places (e.g. The White House) being defined within each.

Environmental Design

Based on the role of spatial knowledge in wayfinding tasks, environmental designers have concerned themselves with developing a design methodology focussed on environmen-

*.We define skilled wayfinding behavior to be purposeful, oriented movement during navigation.

tal organization and map use. Lynch [7] suggests that urban elements such as paths, landmarks, and districts be used to divide the environment into smaller, clearly connected, more manageable pieces. These pieces can then be directly encoded into a hierarchy of spatial knowledge. Lynch also notes the importance of frequent directional cues to orientation maintenance.

Passini [8] expands on these ideas applying them to architectural design. A space should have a basic organizational principle underlying it. For example, Manhattan's streets are organized in a grid. We use this information directly to structure spatial knowledge. Most importantly, a space must have in it a number of "places" which are easily discernible to any wayfinder. A "place" is most simply defined as a distinct, recognizable location or region of a larger space often associated with a landmark. Passini also notes that if a map is to be used, it should show the organizational principle of the space as well as the design elements described by Lynch. The observer's position must always be shown and Levine's forward-up principle must be adhered to.

APPROACH

Previously, we implemented a number of navigation aids and tested their effect on subjects' ability to perform wayfinding tasks in virtual worlds [4]. Although the earlier study addressed the fundamental issues of wayfinding, it did not culminate in a set of generalizable conclusions. Rather, it presented a number of alternative cues and tools which were shown to improve subjects' performance. In this paper we set out to capture subjects' behavior in an objective form from which we could readily infer the search strategies used in accomplishing wayfinding tasks.

Wayfinding Tasks

We classify wayfinding tasks into three primary categories:

1. *Naive search*: Any searching task in which the navigator has no a priori knowledge of the whereabouts of the target in question. A naive search implies that an exhaustive search must be performed.
2. *Primed search*: Any searching task in which the navigator knows the location of the target. The search is non-exhaustive.
3. *Exploration*: Any wayfinding task in which there is no target.

The classifications of wayfinding tasks are mutually exclusive. However, they are often compounded into sequences. In cases where the navigator has general knowledge of the target's position without enough precision to find it directly, a primed search is performed to the target's general proximity followed by a naive search within that area. The opposite ordering of tasks is equally common.

Although purely naive searches are rare in the real world, in virtual worlds, spatial naivete is common in first-time explorers of a space; even the world builder. A scientist visualizing data sets may have no preconceived idea as to the shape or organization of the data. Therefore, wayfinding aids must support both exhaustive and non-exhaustive searches and must facilitate survey knowledge acquisition.

An optimal exhaustive search requires that the navigator traverse the entire space once (in the worst case). To facilitate this, there must be a method of organizing the space to eliminate multiple passes or skipping entire areas. A primed search, on the other hand, requires only that the navigator

know a path to the target. If movement is unrestricted (as it often is in virtual worlds), the navigator need only know the direction and distance to the target. Minimal survey knowledge is required relating the navigator's present position to the target's position. Lastly, exploration is the basic task of spatial comprehension. Its objective is to develop survey knowledge. Maps can be used and, similarly to naive searches, the space should be explicitly organized.

These requirements lead to the conclusion that survey knowledge is the key to successful wayfinding in any environment. Based on the literature previously introduced, we present the following set of design principles for wayfinding augmentations to virtual worlds.

Design Principles

Organizational principles are intended to provide the necessary structure by which an observer can mentally organize the environment into a spatial hierarchy capable of supporting wayfinding tasks. The basic principles are:

1. Divide the large-scale world into distinct small parts, preserving a sense of "place."
2. Organize the small parts under a simple organizational principle.
3. Provide frequent directional cues.

Although map use is not an appropriate tool for every situation, the ability to quickly extract spatial knowledge directly from a map often makes it a powerful navigation aid. Ideally, this knowledge should be flexible, as if the observer had obtained it from direct experience. Therefore, map design principles are intended to present spatial information in such a way as to produce a flexible, orientation-independent representation of the environment. The basic principles are:

1. Show all organizational elements (paths, landmarks, districts, etc.) and the organizational principle.
2. Always show the observer's position.
3. Orient the map with respect to the observer such that the forward-up equivalence principle is accommodated.

Note that the two latter principles are difficult to apply in the real-world when the observer is moving. In the virtual world they can be applied equally easily to a stationary or moving observer.

METHOD

The performance of each subject was observed and measured on wayfinding tasks in several virtual world treatments. Some treatments adhere to a subset of the wayfinding principles; others do not. (See Video Figure in the CHI 96 Video Program) The four treatments in the experiment are:

- **Control Treatment**: No wayfinding assistance provided
- **Grid Treatment**: Adherence to organizational principles
- **Map Treatment**: Adherence to map principles
- **Map/Grid Treatment**: Adherence to both organizational and map principles

In order to capture the subjects' behavior we asked each subject to sketch a map [5, 7] showing the location of land forms and targets after each treatment. We asked subjects to think aloud and recorded each treatment with video and audio tape. Finally, we sampled (once per second) the path of each subject through each treatment and superimposed it

on a map of the environment (See below, Figure 4). In addition to visual comparison to the actual environment, we analyzed the sketch maps using a metric incorporating target distances and directions and land mass shapes developed for this particular purpose [3]. We analyzed the audio and video tape using a verbal protocol analysis [1] and captured observed behavior in the same notation used in CMN-GOMS [2]*.

Design

All subjects were tested on all treatments. Five male and five female subjects participated in the study. The availability of subjects is limited due to both the location of the laboratory and the total time required to complete all trials (approximately three hours for each subject). All subjects had a technical background and were between the ages of 20 and 45.

Stimuli and Apparatus

There were five environments built for this experiment, each of which was equivalent in extent to approximately 12,000 square kilometers in real-world dimensions. The worlds were constructed by hand using a geometric modeling tool. Each world contains large areas of open sea and several land masses. The land masses are colored by elevation and the ocean surface is textured. The land masses were shaped and scaled to be distinct from one another. The targets (ships) were manually placed in the worlds, each with a random orientation. They were also numbered in order to aid identification.

The viewpoint was restricted to movement above the terrain but below a maximum altitude of 400 meters. This is necessary to allow some vertical movement without allowing the subject to gain enough altitude to be able to look down on the entire environment from a "bird's eye view." Movement was constrained horizontally so that when an edge was encountered, there would be no distinguishable features by which to navigate. Contact with the virtual edge resulted in an audible "click" cue combined with immediate clamping of the viewpoint's movement at the point of contact.

The radial grid used in the grid treatment was constructed from a red center post and four different colored posts in each cardinal direction. The outer posts each have a "flag" which points inward toward the red inner post. There are three concentric rings marking range. The red, yellow, and white rings are placed at 10000m, 20000m, and 40000m radius respectively. The outer posts are placed on the white ring. Black radial lines are placed every 45 degrees and extend beyond the white ring.

The map used in the map treatment was identical to the actual environment except that the blue sea texture was replaced by a grey background for contrast with the environment. A red sphere was moved along the map surface to identify the viewpoint position during navigation. The map is presented flat and in the same orientation as the environment itself in conformance to the map principles. The map was placed relative to the viewpoint during movement such that it was visible at all times. The intent was for subjects to feel that the map was in front of their chest.

The map/grid treatment was implemented by placing the grid over the map and the world simultaneously. In the examples in Figure 1, the viewpoint in each case is near the home target facing the center of the world. The target is identified as target zero with a white numbered cube directly above it. The red center post is shown in both the grid and map/grid treatments.

This experiment is intended to study the validity of the design principles and spatial behaviors associated with environmental information rather than the effectiveness of these specific environmental cues. Although we have chosen representative cues for each treatment that adhere to the design principles, there are certainly many other ways to provide the same information.

The treatments were implemented on a Silicon Graphics Onyx™ Reality Engine 2™ workstation. The only peripheral device used was a Fakespace Inc. BOOM3C™ display and tracker. The BOOM3C is a full color, high resolution (1280x1024 pixels) CRT-based display mounted on a counter-balanced mechanical arm. The display is held to the eyes with one hand which controls a single button and a thumb-operated joystick. The position and orientation of the head are tracked through the mechanical arm. Motion is controlled via an acceleration metaphor by which the subject accelerates forward in the virtual world by pushing up on the joystick and backward by pulling down. Velocity is bounded by a maximum speed of approximately Mach 3 (993m/s or 2223mph). Movement is *always* in the direction of view. The subject may stop at any time by pushing the button on the hand controller.

Procedure


The order of treatment presentation and the environments to be used with each were predetermined at random. However, all subjects were given the control treatment on a separate environment as a training session. This was done to introduce the apparatus and the experimental procedure in such a way as to limit the order effect on the data.

The wayfinding task performed for all treatments required the subject to execute five naive searches followed by one primed search. The subject starts at the home target and proceeds to search the environment for each of the five ships which have been numbered and shown to the subject in the instructions. No a priori information is given as to their whereabouts. Once the last target has been located, the subject is required to return to the home target. Subjects were given as much time as needed to complete the task. However, the trial could be discontinued at the subject's request. This was allowed only in cases when either the subject felt unable to make any progress toward task completion after an extended period of time (always at least 15 minutes) or the subject became unable to continue for physical reasons.†


During task execution, subjects were asked to "think aloud" [1] as a method of knowledge elicitation specifically aimed at understanding search strategies. Following each trial, subjects were required to draw a map of the environment in as much detail as possible. Subjects were free to sketch the environment at will starting from a blank piece of paper.

*.Although we used GOMS notation, this is not a traditional GOMS analysis since it was used only for capturing differences in search methods.


†.There were two cases of motion sickness.




The Control Treatment



The Map Treatment



The Grid Treatment



The Map/Grid Treatment

Figure 1: The four treatments as seen from identical locations in a sample environment.

RESULTS AND DISCUSSION

Our primary objective in this paper is to present the fundamental differences in spatial behavior associated with each experimental condition. Consequently, this section will be qualitative in nature. Quantitative results will be reviewed in the next section.

For the purpose of this study, the wayfinding tasks have been broken into their primary components (See Figure 2). Subjects must first ground themselves in the virtual world by acquiring their orientation and position. Then, they begin to undertake the explicit tasks of the five naive searches followed by one primed search. If at any time they lose their orientation or position, they must reacquire it before proceeding. Also, during task execution, actions may be taken which are specifically meant to help the subject maintain orientation and position and consequently, to develop a better cognitive representation. Finally, as targets and objects of interest are located in the world, they are explicitly placed into memory so that they may be recalled later for the map drawing exercise.

This linear description does not completely show the flow of execution. During a single task sequence, subjects may choose more than one method to achieve a goal. For example, when using a word processor, there are typically at least two ways to change the style of the selected text; either via the pull-down menu or using accelerator keys. Either method is acceptable but the subject will never execute both in a single task sequence. In the type of tasks described here, there are also several acceptable methods available to

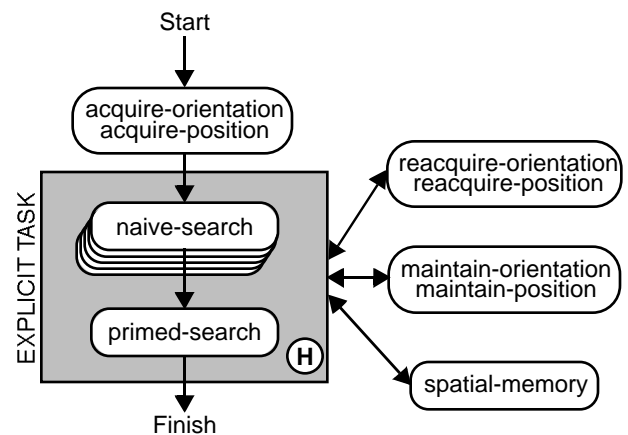


Figure 2: High-level schematic diagram of task execution. The H designates a history mechanism by which any arc traversed out of the gray, explicit task nodes returns to the node of departure.

complete any task, but it *is* feasible to choose more than one method for any given instance. This type of reinforcing or repetitive behavior is the rule rather than the exception in this analysis. Many subjects used several strategies during any one trial as they discovered better alternatives or found their current approach ineffective. Most commonly, subjects selected a method based on intuition restricted by the cues and information currently perceived or available.

The overall goal of the task is to execute an efficient, organized search locating all five targets and afterwards, the home target. The subject must also accumulate survey knowledge in order to draw an accurate map of the environment after completion of the task. Our interest here lies in the notable differences in behavior and in search strategies based on the differences in navigation aids provided (or not provided as the case may be).

Orientation and Position

Grounding the egocentric perspective within the cognitive map, which implies maintaining correct orientation, is an essential component of navigation. Without this grounding, survey knowledge cannot develop. This grounding information must be initially acquired, maintained throughout the task, and reacquired immediately if disorientation occurs.

The lack of suitable orientation cues in the control treatment caused subjects to use relatively weak grounding techniques which proved to be error prone. Most subjects tried to use the orientation of the home target or the initial view direction as a default “north.” This failed, however, when the subject moved away from the start position because adequate reference information was no longer available if and when disorientation occurred. Some subjects tried to use targets and land features as reference points, and attempted to infer distance and direction by moving between them. However, this method was not very successful because of the large distances between targets. One subject even created an external reference by grounding one foot in physical space. The viewpoint was easily reoriented based on the subject’s body posture. However, this method was not useful in extracting positional information. Subjects most often used a dead reckoning* technique with the home target as a reference point to determine position. This was a difficult task considering the lack of alternative reference points throughout the environment. However, it was often the only viable alternative. During navigation, subjects would frequently retrace their recent path or move between targets to estimate distances between them.

These same techniques were also observed in the grid treatment. However, the grid markings were frequently used as the predominant reference points for distance and direction inference. The four outer posts were used as compass points with any one arbitrarily chosen as “north.” Reorientation could be performed by simply returning within view of any outer post. Alternatively, subjects also used the grid markings to divide the space. The colored rings specify range from the center post. Subjects were evenly divided between those who used the posts to divide the space into quadrants and those who used the radial lines to divide the space into wedge-shaped octants. Subjects also used the grid to maintain orientation and position by determining which bounded area they were currently in. These methods were more effective because the grid represented an absolute frame of reference.

The map provides a geocentric perspective to the environment. Consequently, many of the difficulties in acquiring and maintaining orientation and position are eliminated. The map shows the shape of the space and allows for the effective use of geographical landmarks† and world edges as reference points. The orientation of the map itself indi-

*.Dead reckoning is a technique by which future positions are determined based on movement at a constant velocity and direction from a known starting location.

cates the viewpoint orientation with respect to the world. Viewpoint position is available at a glance. The addition of the grid superimposed on the map in the map/grid treatment allows methods from both the map and grid treatments to be used together.

Spatial Memory

This implicit subtask involves the structure of the cognitive map employed by the subject to encode spatial knowledge. Subjects would often pause after locating a target to decide how to remember its location for the map drawing exercise. This is also shown clearly by the subjects’ paths (See Figure 4). Based on analysis of the paths and observed behaviors we defined structures (See Figure 3) which appear to have been used to represent spatial information. The chosen representation was most often closely related to the type of naive search used (See Table 1).

Spatial Memory	Naive Search	Primed Search
Graph	Local Search	Retrace Steps
Square Grid	Lawnmower	Manhattan
Anchor	Coastline	Direct Landmark
Radial Grid	Area Paths	Approximate Direct

Table 1: Spatial memory representations and their associated search methods

The amorphous nature of the control treatment allowed subjects to apply any structure to the environment. The targets could be connected in the form of a *graph* with the targets as nodes and the paths between them as edges. A *square grid* structure could also be applied dividing the space into smaller regions. Most commonly, targets were represented as *anchored* relative to the land masses.

In the grid treatment, the graph and anchoring methods are still utilized. However, the square grid method does not coincide with the shape of the radial grid and is consequently replaced by a *radial grid* structure.

The map does not impose any structure on the environment beyond that of the control treatment. Accordingly, the same conceptual structures were used in the map treatment as in the control treatment. In the map/grid treatment, however, subjects were not observed to use the graph method; presumably because the grid markings and geocentric perspective provided enough reference points to make this method inferior to others.

Naive Searching

An effective naive search requires that the subject conduct an organized, exhaustive traversal of the space. In the control treatment, since the space itself was not organized, subjects had extreme difficulty in organizing their search. The most common methods involved searching the area *local* to targets or to land masses by following *coastlines*. The local search method was most often used in concert with the graph structure while the coastline method was used with the anchor structure. The coastline method is improved by

†.These include unusual landforms such as a bay or peninsula.




Figure 3: Cognitive representations for spatial memory

encircling the land multiple times, each time extending further out into open space. The subject using an external orientation method was able to conduct a very different type of search because of his superior ability to maintain orientation. He located a corner of the world by repeatedly bumping the edges, and then proceeded to cover the space in long strips moving across the length of the world (not unlike the path of a lawnmower; See Figure 4A)

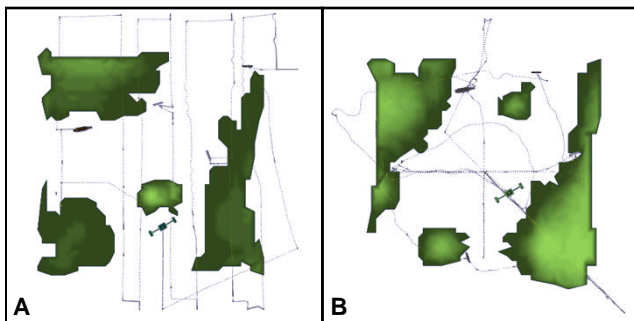


Figure 4: Paths of subjects with: A. The lawnmower naive search method. B. The radial grid used as paths.

When the grid was provided, it added two primary methods to those of the control treatment. The grid markings can be used as *paths*. The path trail line in Figure 4B clearly shows both the radial lines and the range rings which were used to guide the search. This was seen as a more comfortable alternative to searching open spaces. Also, the world could be divided into *areas* bounded by the grid lines, each to be searched sequentially.

The map allowed the search techniques of the control treatment to be executed more effectively and with optimizations. *Heuristic* searches were often used to initially visit places which subjects guessed would hide a target such as an inlet or bay. They can be used with any of the spatial memory representations. Also, the lawnmower method was optimized to eliminate searching over land. The addition of

the grid to the map enabled combinations of methods. For example, heuristics were used to improve area searches.

Primed Searching

The ability to perform an efficient primed search is dependent on the level of useful spatial knowledge obtained during navigation. In the control treatment, many subjects had been disoriented for most of the trial and therefore conducted another naive search to locate the home target. Subjects who were able to connect targets in a chain-like fashion *retraced steps* back to familiar ground. If a grid structure was used to organize the space, a path was followed along the *Manhattan distance** to the home target. Most often, subjects unsuccessfully attempted to infer a *direct* path.

The grid markings were used in one of two ways. If a post or line crossing was within view of the home target, that location was used as an *approximate* target from which the primed search could be completed easily. The alternative was to place the home target in the world relative to the grid and search for it directly.

When the map is used, the home target is placed relative to some predetermined geographical *landmark*. To relocate the target, the subject moves back to the general area and conducts a local search until successful. If the grid is added to the map, the grid markings are usually preferred over geographical landmarks for reference points.

In all treatments, dead reckoning was used extensively as a way to structure movement during a search (naive or primed). It was used frequently in the control and grid treatments but only intermittently in the map and map/grid treatments when not navigating from the map directly.

SEARCH EFFECTIVENESS

In addition to categorizing searching strategies and behavior, we evaluated search effectiveness across treatments. A Friedman two-way within-subject ANOVA was used to analyze the results. This is a nonparametric test of interval data suitable for experiments with small subject sample sizes.

A measure of total time to complete each trial (adjusted for incomplete attempts) showed significant effects across treatments (Friedman test statistic=15.96, p .001)[†] indicating that subjects in the control treatment tended to execute ineffective searches and were unable to structure the space. This is further illustrated in the distance travelled during each treatment which also showed a significant effect (Friedman test statistic=17.4, p .001). As seen in Figure 5, subjects repeatedly traversed the same space in the control treatment. Distances and times for the grid treatment were lower than the control treatment but higher than the map and map/grid treatments. Although the grid supplies enough information to structure a search, a significant amount of effort and reinforcing actions must be executed to maintain orientation.

A ratio of the percentage of space searched to total time is a coarse indicator of search effectiveness. This ratio was shown to be significant across treatments (Friedman test statistic=13.8, p .005). In general, subjects were not able to coordinate an effective search in the control treatment due to

*.A Manhattan distance measures movement constrained to horizontal and vertical directions.

†.For all quantitative results presented, no significant effect was found for gender or treatment order.

disorientation. The grid significantly increased performance and ability to extract spatial information but the map and map/grid treatments supported the most effective searches.

A metric was developed to objectively evaluate the map sketches (See [3] for details) in terms of directional accuracy, relative distance estimation, and land form (shape, placement, and relative scale of the land masses). There were significant effects for directional accuracy across treatments (Friedman test statistic=14.52, p .005) as well as land form (Friedman test statistic=18.84, p .0001). However, distance estimation did not show a significant effect (Friedman test statistic=6.93, p >.05). Error in directional accuracy (range from 3%-60%) was significantly higher than error in distance estimation (range from 1%-10%). The ability of most subjects to efficiently dead reckon regardless of treatment is dependent on their ability to determine distance and speed. Therefore, it is not surprising that subjects were much more adept at estimating distance than direction in light of the prevalence of dead reckoning across treatments. Furthermore, disorientation often caused subjects to draw incomplete map sketches of the control treatment because they were unable to do more than randomly place targets on the map which were located while they were disoriented.

The disorientation most commonly associated with the control treatment is evident in a number of ways. A typical diagram of the route followed during an average control treatment shows significant confusion and a general inability to organize the space (See Figure 5). In this example, a




Figure 5: A typical control treatment path and its associated map drawing. Targets have been numbered on both figures.

coastline following method was used. The weaving motion during navigation is a behavior associated with widening the field of view. The subject looks back and forth while moving to compensate for the narrow view of the display. The path shows that this subject traversed each land mass multiple times with a single thread joining A to the others. There is no point in this environment from which A can be seen simultaneously with B or C. Consequently, most of the search was confined to familiar areas near the home target (labelled #0). This subject drew a map which shows a basic inability to extract spatial information from the space. The map is roughly a horizontal reflection of the actual environment.

The presence of the radial grid was shown to significantly improve directional accuracy. Map sketches of both the grid and map/grid treatments were more accurate in target placement than the control and map treatments. This is shown in Figure 6. This subject used an area search method in the grid treatment. The map shows that although the subject

placed the targets very accurately, the land mass structure is incorrect. Due to the geocentric view provided, the map and map/grid treatments showed the most accurate land forms on the map sketching exercise.

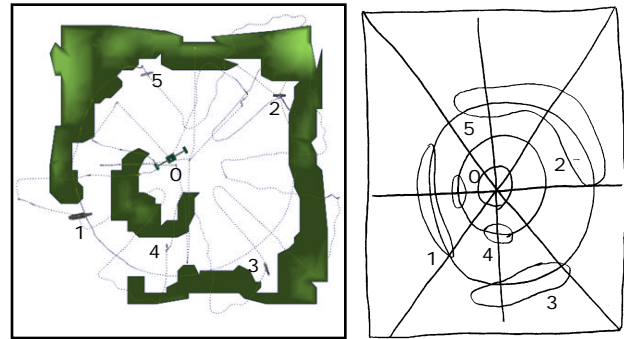


Figure 6: A grid treatment showing the patterns of search confined to sequential areas of the grid.

The fact that we were able to discern reasonable cognitive representations (Figure 3), even in the control treatment indicates that structure was imposed on the environments whether or not it was supported by the environment itself. This structure is a necessary precondition to execution of a searching task. Those subjects who did not conceptually organize the environment were unable to conduct an efficient search and in many cases, did not successfully complete the task.

Thorndyke [10] describes survey knowledge in part as configurational knowledge, often acquired via a map, allowing short-cuts and the ability to infer new paths. The optimizations to the control treatment search methods observed in the map treatment illustrate this point exactly. The path shown in Figure 7 is cut short around target 5 rather than cross a land mass of known shape. This path can be compared to Figure 4A which shows the same search technique used in the control treatment. The subject is unsure of the shape and size of the land masses and therefore does not deviate from the pattern.




Figure 7: A map treatment illustrating the lawnmower method of search.

Behavior observed during all four treatments indicates that organizational structure and restriction of movement are an essential part of navigation. Subjects were uncomfortable with completely free and unrestricted movement in sparsely populated space because it led to disorientation much of the time. When the grid was added, this discomfort subsided as subjects found structure to guide and direct their search and consequently, avoid disorientation. The importance of paths

as described by Lynch [7] indicates that in virtual and physical spaces alike, navigators need structured movement for effective navigation and spatial knowledge acquisition. This should not be interpreted as a recommendation of rigid movement restriction in virtual worlds. A lack of freedom in movement and choice inhibits exploration. Rather, we advocate the augmentation of large virtual worlds with directional cues and a simple, well-defined structure that facilitates subjects' construction of their own paths.

The frequent use of dead reckoning across treatments was not an anticipated behavior. Effective use of this technique requires three necessary components:

- the current position
- the direction of movement, and
- the velocity of movement.

The effectiveness of the technique is much improved if there are reference points within view of the path of travel which can be used as "checkpoints" marking progress along the way. This was only the case when the grid was present. From the relative success of dead reckoning in the non-grid treatments we conclude that the frame rate (15 frames/second) and visual cues (specifically the sea texture) provided enough feedback to maintain a consistent estimation of velocity. The technique is further encouraged by the method of movement since the subject must always look in the direction of movement.

CONCLUSIONS

The methods employed in each subtask were profoundly affected by the stimuli presented. The lack of directional cues and spatial organization in the control treatment led to ineffective search strategies and frequent disorientation. The radial grid provided enough information to successfully execute a search but required reinforcing actions to maintain orientation. The map provides a simultaneous geocentric perspective augmenting the egocentric perspective and fostering the use of geographical landmarks and optimizations to search methods. General conclusions drawn from this work include:

1. When not given an adequate source of directional cues, disorientation will inhibit both wayfinding performance and spatial knowledge acquisition.
2. A large world with no explicit structure is difficult, if not impossible, to search exhaustively. This was shown by repeated reacquisition behavior in the control treatment.
3. A conceptual coordinate system is often imposed on the world to act as a divider. This is a side-effect of not being able to divide the world explicitly. A structure must be imposed on the world if an organized exhaustive search is to be attempted.
4. Our observations support the notion that path following is a natural spatial behavior. Subjects frequently used features such as coastlines or grid lines as if they were paths.
5. A map allows for optimizations to be made to search strategies. This is because it can be considered a supplement to survey knowledge.
6. Dead reckoning was observed to be an intuitive and nat-

ural part of navigation; all subjects exhibited this behavior even though frequently unaware of it. The ability to infer position from a past location and constant velocity over time, while sometimes complex in reality, appears to be more easily understood and implemented in virtual spaces.

We are encouraged that the observations reported in this paper support the application of environmental design principles to virtual world design. We expect to broaden our investigation in the future to include environments with vastly different spatial characteristics. As environments become more complex and abstract, we will need to determine whether or not a human's conception of an abstract space is analogous to that of a physical space. Early indications are that this is the case, but a definitive statement at this time requires us to speculate as to the nature of virtual environment representations of the future.

ACKNOWLEDGEMENTS

This research was conducted through support by the Virtual Environment Laboratory of the Tactical Electronic Warfare Division of the Naval Research Laboratory in Washington, D.C. The authors would like to thank Manuel Pérez and Dan McFarlane for their assistance as well as our experimental subjects at NRL for their time and effort.

REFERENCES

1. Byrne, R., *Protocol Analysis in Problem Solving, in Thinking and Reasoning: Psychological Approaches*, J.B.T. Evans, Editor. 1983, Routledge & Kegan Paul: London. p. 227-249.
2. Card, S.K., T.P. Moran, and A. Newell, *The Psychology of Human-Computer Interaction*. 1983, Hillsdale, N.J.: Lawrence Erlbaum Associates.
3. Darken, R.P., *Wayfinding in Large-Scale Virtual Worlds*, Doctoral dissertation, 1996, The George Washington University, Washington, D.C.
4. Darken, R.P. and J.L. Sibert, *A Toolset for Navigation in Virtual Environments*. User Interface Software and Technology, 1993: p. 157-165.
5. Howard, J.H. and S.M. Kerst, *Memory and Perception of Cartographic Information for Familiar and Unfamiliar Environments*. Human Factors, 1981. **23**(4): p. 495-504.
6. Levine, M., I. Marchon, and G. Hanley, *The Placement and Misplacement of You-Are-Here Maps*. Environment and Behavior, 1984. **16**(2): p. 139-157.
7. Lynch, K., *The Image of the City*. 1960, Cambridge: MIT Press. 194.
8. Passini, R., *Wayfinding in Architecture*. 1984, New York: Van Nostrand Reinhold Company Inc.
9. Stevens, A. and P. Coupe, *Distortions in Judged Spatial Relations*. Cognitive Psychology, 1978. **10**: p. 422-437.
10. Thorndyke, P.W. and S.E. Goldin, *Spatial Learning and Reasoning Skill*, in *Spatial Orientation: Theory, Research, and Application*, H.L. Pick and L.P. Acredolo, Editors. 1983, Plenum Press: New York. p. 195-217.